

METHOD AND SYSTEM FOR VERIFYING THE INTEGRITY OF A DATABASE

Field of Invention

5 The present invention relates to a method and system for verifying the integrity of a database for use by an application. More particularly, but not exclusively, the present invention relates to a method and system to verify the integrity of a database for use by an application by comparing schema metadata extracted from the database with previously extracted schema metadata.

10

Background to the Invention

Where a central database is servicing multiple applications, either over a network or on a single computer system, and an application has the capacity to change the structure of the database, there is often a need for the other applications to monitor such changes to ensure that any changes made do not affect the operation of those applications.

Presently, applications do not check the structure of a database to assess whether it remains compatible with the application. An application is often only aware of critical changes to the structure of the database when a database query made by the application fails.

In order for the application to utilise a database effectively it is desired that the application is assured of the integrity of the structure of database before using it. Data which describes the structure of the database is called schema metadata. Schema metadata includes tables, columns in tables, datatypes of columns, lengths of columns, custom database data types, foreign keys, constraints, stored procedures, views, triggers, indices, and scheduled jobs.

30 To an application the variation of some elements of schema metadata may not be important. Therefore what is required is a method to verify the integrity of a database for use by a specific application. Integrity of the database concerns structural changes to the database which affect that particular application. Thus where structural changes do not affect an application the integrity of the database is maintained.

It is an object of the invention to provide a method and system to address the above problems or to at least provide the public with a useful choice.

5 **Summary of the Invention**

According to a first aspect of the invention there is provided a method for a primary application to verify the integrity of a secondary application including the steps of:

- i. obtaining a reference reduced representation by:
 - 10 A. applying a process to obtain schema metadata from the secondary application;
 - B. creating a reference reduced representation of the first obtained schema metadata using an algorithm; and
 - C. storing the reference reduced representation;
- 15 ii. during execution of the primary application, applying the process to obtain the schema metadata from the secondary application;
- iii. creating a second reduced representation of the second obtained schema metadata using the algorithm;
- iv. comparing the reference reduced representation and the second reduced representation; and
- 20 v. controlling execution of the primary application dependent on the outcome of the comparison.

In a preferred implementation of the method the secondary application is a database.

25

The schema metadata may include tables, columns in tables, datatypes of columns, lengths of columns, custom database data types, foreign keys, constraints, stored procedures, views, triggers, indecies, or scheduled jobs.

30 Preferably the algorithm to create the reduced representations is a hash function, consequently the reduced representation is a hash of the schema metadata. The hash function may be MD5, CRC32, or any other contemporary hashing algorithm.

The algorithm may be a lossless compression algorithm, such as zip, gzip, or bzip2.

The reference reduced representation may be embedded within the source code for the primary application or within a configuration file for the primary application.

5 Step (i) for obtaining the reference reduced representation may be repeated more than once before steps (ii) to (v). This preferably occurs when an expected change occurs to the schema metadata of the database and a new reference for the database is required by the primary application.

10 Preferably the process includes organising the extracted schema metadata using a nested and determinable method. The method may be by alphabetical listing, default database order, creation date order, or by table owner.

15 In step (v) execution of the primary application may be controlled by halting the primary application, or sending an error message to any of the user, the application or database manager, or to the database itself.

A schema stability lock may be requested of the database. A schema stability lock prevents changes to the database schema. The lock may be obtained before step (iii).

20 The lock may be maintained if the comparison is successful (the reference reduced representation and the second reduced representation match) in order that application can utilise the database without fear of the schema changing. The lock may be released if the comparison is unsuccessful. Alternatively, the lock may be obtained after step (iv) if the comparison is successful.

25 The process may be defined so that it obtains all available schema metadata. The process may be defined so that it only obtains schema metadata relevant to the operations of the primary application.

30 The process may utilize SQL92 standard which is widely implemented by databases such as Oracle and SQL Server. An example of this is the view "INFORMATION_SCHEMA.TABLES" which lists the tables in a database.

The process may utilize the services of a standard driver API for the database. The

process may utilize the JDBC API provided within Java. Contained in JDBC is a class called "java.sql.DatabaseMetaData" which provides mechanisms to get tables in a schema and other metadata from the database.

- 5 According to a further aspect of the invention there is provided a system for verifying for a plurality of applications the integrity of one or more databases including:
- i. a plurality of applications adapted to store a plurality of previously calculated reduced representations of schema metadata for one or more databases, to extract a plurality of schema metadata from one or more
10 databases, to newly calculate a plurality of reduced representations from the plurality of extracted schema metadata, and to compare each of plurality of previously calculated reduced representations with its corresponding newly calculated reduced representation; and
 - ii. one or more databases adapted to receive requests for schema metadata
15 from the plurality of applications and to transmit schema metadata to the plurality of applications.

The plurality of applications may be deployed on a plurality of user computer devices. The one or more of databases may be deployed on one or more server computer
20 devices. The user computer devices and the server computer devices may be connected via a network such as LAN, WAN, or the Internet.

- According to a further aspect of the invention there is provided a system for verifying for an application the integrity of a database including:
- 25 i. an application;
 - ii. a stored reduced representation of schema metadata of a database; and
 - iii. a verification engine which upon connection to a database obtains a reduced representation of schema metadata from the database and compares it with the stored reduced representation to control the
30 application.

Brief Description of the Drawings

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1: shows a flow diagram illustrating the method of the invention.

Figure 2: shows a flow diagram illustrating the verification stage of the method for a plurality of databases and applications.

Figure 3: shows a sequence diagram illustrating the reference creation stage of a first example of the method.

Figure 4: shows a sequence diagram illustrating the verification stage of a first example of the method.

Figure 5: shows a sequence diagram illustrating the reference creation stage of a second example of the method.

Figure 6: shows a sequence diagram illustrating the verification stage of a second example of the method.

Figure 7: shows a diagram illustrating how the method may be deployed on hardware.

Figure 8: shows a diagram illustrating a possible implementation of the method.

Detailed Description of Preferred Embodiments

The present invention relates to a method and system for verifying the integrity of a database for use by an application.

The invention will now be described with reference to Figure 1 of the drawings.

Reference Creation Stage

During the build (creation) of an application, a database 1, which the application will during its execution, is queried to extract in step 2 schema metadata 3. Schema metadata is information that describes the structure and other features of a database and is agnostic to the actual data stored in the database. Schema metadata includes
5 tables, columns in tables, datatypes of columns, lengths of columns, custom database data types, foreign keys, constraints, stored procedures, views, triggers, indices, and scheduled jobs.

The set of schema metadata that is extracted is typically that metadata which would
10 adversely affect operation of the application if it changed. Therefore it may not be necessary to extract the entire set of schema metadata. Circumstances such as an application's sensitivity to indices will vary from application to application and the set of schema metadata that is to be extracted shall vary appropriately.

15 For some applications full schema extraction will be best. Full schema extraction is particularly desirable for many large and performance orientated applications. Full schema extraction may include not only the schema in which data resides but also other aspects such as indices if these are essential to the correct functioning of the application.

20 For applications that are installed in multiple places, it may be desirable to ignore the indices that are present in the database. For example two sites may run the same application, but use that application in different ways. For performance reasons the sites may choose to have different indices to reflect the nature of their data and usage.

25 Some applications allow user defined tables. In such circumstances it is not possible to create a reference reduced representation for as yet undefined tables when the application is built. In such cases only tables that will be constant for the application should be included in the set of schema metadata to be extracted.

30 For example, if an application only needs the database to query an "email contacts" table to extract a "phone number" by "name", then the metadata the application is concerned about is (a) the existence of a "email contacts" table, (b) the existence of a "name" column, and (c) the existence of a "phone number" column. The application is

unconcerned about the existence of an “email addresses” column. Correspondingly, the metadata extracted by the application would be (a), (b), and (c).

There are a number of ways to extract schema information.

5

Database vendors generally have product and version specific ways to extract data from their database. SQL Server has version specific stored procedures such as “sp_tables” for listing tables in a database and “sp_columns” for listing columns in a table.

10

There is a vendor neutral method for retrieving schema information that is part of the SQL92 standard and widely implemented by databases such as Oracle and SQL Server. An example of this is the view “INFORMATION_SCHEMA.TABLES” which lists the tables in a database.

15

Another method for getting schema information is to use the services of a standard driver API. Java has a specification for dealing with a database called JDBC. Contained in this is a class called “java.sql.DatabaseMetaData”. This provides mechanisms to get tables in a schema and other metadata from the database.

20

The extracted metadata is organised using a method which results in a nested and determinable order, such as alphabetically. For example, if the schema metadata to be extracted was a “Clients” table with columns “Name”, “Address”, “Phone” and a “Files” table with columns “File Number”, “Client”; then the resulting ordered metadata list would be “Clients, Address, Name, Phone, Files, Client, File Number”.

25

Other methods of ordering include:

- by default database order
- by creation date order
- by table owner

30

The organised extracted metadata is compressed in step 4 using a hash function into a reduced form 5. A good hash function will produce a difficult to forge representation which uniquely identifies the schema metadata. Examples of good contemporary hash functions include MD5 and CRC32.

It will be appreciated that other methods of cryptography or compression may be used to create the reduced form of the metadata, including, lossless compression algorithms such as zip, gzip, and bzip2.

5

This reduced form or hash is stored in step 6 as a reference for the correct configuration of the database. Preferably, the reduced form is stored by embedding into the application or within an application configuration file.

10

The reference creation stage may occur again during the life of the application. For example, when the application is aware (or initiates) an expected change to the structure of the database.

Verification Stage

15

During execution of the application, the database 1 is queried to extract in step 7 the same type of schema metadata 8 as during the Reference Creation Stage. This metadata is organised using the same method used during the Reference Creation Stage and then compressed in step 9 using the same algorithm used during the Reference Creation Stage to create a runtime reduced form 10. The runtime reduced form 10 is compared in step 11 with the reference reduced form 5 which was embedded in the application.

25

The outcome 12 of this comparison controls the execution of the application.

30

For instance, when the runtime reduced form differs from the embedded reduced form, this indicates that relevant metadata within the database has changed. If the runtime reduced form is identical to the embedded reduced form this indicates that it is very unlikely that relevant metadata within the database has changed. In the case of the former, the application may stop execution, generate an error message, or take actions against the database. In the case of the latter, the application can confidently query the database in the knowledge that the queries are likely to succeed.

The verification process may occur several times during the execution of the application,

especially in the case where the application continues execution over a significant period of time, or it may only occur at the start of the application. To ensure that the database does not change after the verification process and before a query is made of the database, a schema stability lock, such as a "SQL Server 2000 schema stability lock",
5 may be requested of the database.

A SQL Server 2000 schema stability lock on a table prevents other processes from acquiring locks that alter the structure of a table. The abbreviated name for this lock type is "Sch-S". The method may use schema stability locks to guarantee that the
10 schema of tables doesn't change between database integrity verification. This prevents the application getting out of match with the database even if some process is attempting to alter the database structure.

Referring to Figure 2, a second implementation of the invention will be described.

15 In this implementation there are a plurality of databases 13, 14, and 15 and a plurality of applications 16, 17, and 18. Several of the databases 14 and 15 are utilised by more than one application, and several of the applications 16 and 18 utilise more than one database.

20 For each of the applications, reduced representations of each utilised database's metadata has been stored in a manner accessible by that application. The representations can be stored in the application's configuration file, or embedded within the application itself if the representations were obtained at or before build time of the
25 application.

Figure 2 shows the verification process used by each application to ascertain the integrity of the databases they are using.

30 Application 16 is using databases 13 and 15. Therefore during execution of application 16 the following will take place:

- i) Metadata 19 referring to aspects of database 13 used by application 16 is extracted from database 13, a hash algorithm is applied in step 21 to form a

reduced representation 22 of the metadata and this is compared in step 23 with the stored reduced representation 24 to confirm the integrity of database 13 for use by application 16.

- ii) Metadata 25 referring to aspects of the database 15 used by application 16 is extracted in step 26 from database 15, a hash algorithm is applied in step 27 to form a reduced representation 28 of the metadata and this is compared in step 29 with the stored reduced representation 30 to confirm the integrity of database 15 for use by application 16.

Application 17 is using database 14. Therefore during execution of application 17 the following will take place:

- i) Metadata 31 referring to aspects of database 14 used by application 17 is extracted in step 32 from database 14, a hash algorithm is applied in step 33 to form a reduced representation 34 of the metadata and this is compared in step 35 with the stored reduced representation 36 to confirm the integrity of database 14 for use by application 17.

Application 18 is using databases 14 and 15. Therefore during execution of application 18 the following will take place:

- i) Metadata 37 referring to aspects of database 14 used by application 18 is extracted in step 38 from database 14, a hash algorithm is applied in step 39 to form a reduced representation 40 of the metadata and this is compared in step 41 with the stored reduced representation 42 to confirm the integrity of database 14 for use by application 18.
- ii) Metadata 43 referring to aspects of the database 15 used by application 18 is extracted in step 44 from database 15, a hash algorithm is applied in step 45 to form a reduced representation 46 of the metadata and this is compared in step 47 with the stored reduced representation 48 to confirm the integrity of database 15 for use by application 18.

The stored reduced representation 36 for database 14 for application 17 is different to the stored reduced representation 42 for database 14 for application 18 because

application 17 and application 18 use different aspects of database 14.

The outcomes 49, 50, 51, 52, and 53 of the comparisons 23, 29, 35, 41, and 47 will control the execution of each application.

5

It will be appreciated that any lossy or lossless algorithm which creates reduced representations may be used. There is no requirement to recreate the actual metadata extracted therefore the preferred algorithm favours high compression over data loss.

10

There is a requirement that any minor changes to the metadata are detected therefore an algorithm with properties similar to a hash function is preferred.

Example 1

15

A preferred implementation of the invention will now be described with reference to Figures 3 and 4.

Reference Creation Stage

20

In this example an application developer 54 instructs in step 55 the build script 56 to build the application 57. The build script 56 calls in step 58 a function from a module "schema utilities" 59 which obtains in step 60 a schema from RDBMS1 61 – the data source which is to be used by the application – and calculates a fingerprint – a reduced representation of the schema which can be used to identify the schema in future – from the schema data

25

The fingerprint is embedded in step 62 in the application 57.

Verification Stage

30

A user 63 of the application 57 starts in step 64 the application. Sometime during its execution, the application 57 calls in step 65 a function from the module "schema utilities" 59 which obtains in step 66 a schema from RDBMS1 61 and calculates a new fingerprint from the schema.

The new fingerprint is compared in step 67 with the embedded fingerprint.

If the RDBMS1 schema has not changed since the application was built then the comparison will return a match. If the fingerprints do not match then the schema has changed. The application may choose to stop execution at this point, it may notify the user or a database/application manager, or it may send a message to the data source.

Example 2

An implementation of the invention will now be described with reference to Figures 5 and 6. This implementation relates to an application which needs to verify the integrity of two data sources.

Reference Creation Stage

In this example an application developer 68 instructs in step 69 a build script 70 to build the application 71. The build script 70 calls in step 72 a function from a module "schema utilities" 73 which obtains in step 74 a schema from RDBMS1 75, a first data source to be used by the application, and calculates a fingerprint from the RDBMS1 schema data.

The build script 70 calls in step 76 a function from a module "schema utilities" 73 which obtains in step 77 a schema from RDBMS2 78, a second data source to be used by the application 71, and calculates a fingerprint from the RDBMS2 schema data.

Both fingerprints are embedded in step 79 in the application.

Verification Stage

A user 80 of the application starts the application 71. Sometime during its execution, the application 71 calls in step 81 a function from a module "schema utilities" 73 which obtains in step 82 a schema from RDBMS1 75 and calculates a new fingerprint from the schema.

The new fingerprint is compared in step 83 with the embedded RDBMS1 fingerprint.

If the RDBMS1 schema has not changed since the application was built then the comparison will return a match. If the fingerprints do not match then the schema has changed. The application may choose to stop execution at this point, it may notify the user or a database/application manager, or it may send a message to the data source.

Sometime during its execution, the application 71 calls in step 84 a function from a module "schema utilities" 73 which obtains in step 85 a schema from RDBMS2 78 and calculates a new fingerprint from the schema.

The new fingerprint is compared in step 86 with the embedded RDBMS2 fingerprint.

If the RDBMS2 schema has not changed since the application was built then the comparison will return a match. If the fingerprints do not match then the schema has changed. The application may choose to stop execution at this point, it may notify the user or a database/application manager, or it may send a message to the data source.

Figure 7 illustrates how the invention may be deployed on hardware.

The primary application and secondary application are deployed over a network 87. The network may be a LAN, WAN, or the Internet.

The secondary application is typically a database and in this example is deployed on a server 88 linked to the network 87. The primary application is typically a user application and in the example is deployed on a personal computer 89 linked to the network 87.

It will be appreciated that the personal computer may be of any hardware configuration, such as an IBM Pentium 4 with 256 MB of RAM and a 40 GB hard drive.

In situations where there is a plurality of secondary applications these may be deployed on separate servers or on the same server. In this example an additional database is deployed on a separate server 90.

In situations where there is a plurality of primary applications these may be deployed on

separate personal computers or on the same personal computer. In this example an additional primary application is deployed on a separate personal computer 91.

It will be appreciated by a skilled reader that both the primary application and the secondary application may be deployed on a single computer. For example, a single IBM Pentium 4 with 256 MB of RAM and a 40 GB hard drive running a Windows 2000 operating system

It will further be appreciated that a plurality of primary applications and secondary applications may be deployed on a single computer. For example, a single IBM Pentium 4 with 256 MB of RAM and a 40 GB hard drive running a Windows 2000 operating system, including an ORACLE database for a payroll (first database) and a MICROSOFT ACCESS database (second database) for office stationary. The computer further including a first primary application, such as an integrated office management application interfaced to both databases, and a second primary application, such as an office stationary monitoring application interfaced to the second database.

Referring to Figure 8 an example of how the invention might be implemented will be described.

The example concerns a client/server architecture timekeeping system for professionals, such as lawyers. Such a system is comprised of a central database 92 storing all data, and a number of client applications 93, 94 and 95 on the professionals' desktops 96, 97 and 98.

The client applications 93, 94, and 95 communicate with the database server 99 directly. The database 92 is a SQL Server 2000 database. The client applications 93, 94, and 95 are written in Java.

Version 1.1 of the client application was initially deployed on all the desktops. This version had a stored hash value 100 of schema metadata extracted from the database.

The timekeeping system was modified and changes were made to the schema of the database. The new 1.2 version of the client application 93 and 94 was rolled out to the

desktops 96 and 97. The version 1.2 client applications 93 and 94 have hash values 101 and 102 of the version 1.2 database's modified schema metadata embedded within them.

- 5 During the upgrade process, one desktop 98 missed out on the new version and is still running version 1.1 of the client application 95.

When the version 1.2 client applications 93 and 94 on the desktops 96 and 97 they extract schema metadata in steps 103 and 104 from the database 92, compute in steps
10 105 and 106 the hash values 107 and 108 of the schema metadata, and compare in steps 109 and 110 the computed hash values 107 and 108 to the stored hash values 101 and 102. As the database is also version 1.2, the comparison reveals a match between the computed hash values 107 and 108 and the stored hash values 101 and 102. As a result of the match the applications continue normal execution, in steps 111
15 and 112, which involves use, in steps 113 and 114, of the database.

When version 1.1 of the client application 95 executes, on the desktop 98 that missed out on the upgrade, it extracts schema metadata 115, computes 116 the hash value, and compares 117 it with the stored hash value 100. As the stored hash value 100 is for
20 version 1.1 of the database, the comparison will return a mismatch. As a result of the mismatch the client application will cease execution 118 in order to avoid potential data loss within the database 92.

The invention provides numerous advantages including:

- 25
- potential elimination of schema and SQL statement mismatches resulting in runtime errors;
 - potential to reduce chance of data loss due to schema mismatch;
 - detection of mismatches of various kinds before any data processing is attempted;
- 30
- various schema matching algorithms to detect only the kind of schema mismatches that an application is sensitive to;
 - storing the schema in a reduced form to minimise storage space in client applications;
 - by schema matching on meta data instead of literal data such as a table with an

application version number, the possibility that the literal data could be erroneously modified is eliminated; and

- detection of mismatch at the earliest possible time.

5 While the present invention has been illustrated by the description of the embodiments thereof, and while the embodiments have been described in considerable detail, it is not the intention of the applicant to restrict or in any way limit the scope of the appended claims to such detail. Additional advantages and modifications will readily appear to those skilled in the art. Therefore, the invention in its broader aspects is not limited to the
10 specific details representative apparatus and method, and illustrative examples shown and described. Accordingly, departures may be made from such details without departure from the spirit or scope of applicant's general inventive concept.